License: Creative Commons BY-NC-SA 4.0
This document was written by **Anatole Ternaux** and may be shared and adapted for non-commercial use with proper attribution and under the same license. **https://creativecommons.org/licenses/by-nc-sa/4.0/**

29 June 2025

Welcome to the Churn: Generative AI, What it Is, and What it Isn't, and How You Can Fight Back

Originally published: https://durocsdeskdrawer.ink/2025/06/29/welcome-to-the-churn-generative-ai-what-it-can-do-and-what-it-cant-and-how-you-can-fight-back/

You've seen it, haven't you? Fantasy books that are chasing the next ACOTAR high, cookie-cutter vaguely-medievalish settings that have the depth of vermeil over cheap steel, where the draw is the tropes and the vibes, and not the world.

Fantasy, YA fantasy, romantasy, YA dystopia, romance in general — these are genres that are a bloating corpse left out in a swamp, and the maggots feasting on it are AI-generated.

This speaks for itself

What's that about AI-generation? You see, the modern crop of those genres follow tropes, they follow a template. "Save the Cat," where the protagonist comes with built in trauma, discover their powers through healing sex, and have most of the messy threads wrapped up neatly at 70% of the book. AI writes like fantasy, and fantasy writes like AI.

This is how you get controversies like *Silver Elite*, which many readers have charged as being AI-generated slop due to its structure and lack of an authorial presence. And also how you get Lena McDonald and her *Darkhollow Academy* scandal, where McDonald left a ChatGPT prompt in the final version of her second book.

K.C. Crowne was another one caught leaving AI prompts in her book. Rania Faris too. Tim Boucher shamelessly posts "his" fantasy books on Youtube, showing viewers how he pushed out 120 books in two years with AI and how to game Amazon KDP.

Stephen Marche (under pen-name Aidan Marchine) released the novella Death of an Author, confessing that 95% of it was AI-generated through ChatGPT, Cohere, and Sudowrite. It's not straight up fantasy, but sci-fi with fantasy components.

And aside from that last example, do you see what each and every one of these AI cranks have in common? They write fantasy, and fantasy taught the machines paint-by-numbers plot mechanics.

Templates and tropes are easy for AI to mimic and follow, contributing to the bloat. I recently heard a statistic of literary agents being flooded with 60-70% fantasy books, many of which read suspiciously AI-generated. It's gotten to the point where publishing houses and literary agents turn down fantasy submissions without even bothering to read them.

Soon, fantasy and fantasy-adjacent genres will be like a glam metal band in 1994: unwanted, unwelcome, and no one is listening because all the kids are into grunge instead. Fantasy will survive in some form, but not in its current algorithmic churn. Rather, it will survive by going back to literate, weird fantasy, like Moorcock's *Elric* series, Beagle's *The Last Unicorn*, Duane's *Young Wizards* series, and L'Engle's *A Wrinkle in Time*.

Likewise, savvy and literate readers are noticing the templates, and are experiencing increasing trope and genre fatigue. It's always the same kind of courts, the same kind of protagonist, the same kind of execution that hits that particular dopamine high. The AI templates make it easy to churn out more. But what happens when that dopamine wears off, and you need more, and more and more to maintain the high? Those readers either dig in and demand more, or they leave, looking for something fresher, something clearly not made by a machine.

As they say in Corey's *The Expanse*, "**Welcome to the Churn**." This is where you either float to the top, or sink to the bottom, and everything in-between is the Churn. The Churn is what happens when there is major upheaval and things do not go back to the way they were. It's when you have to break old habits, when norms shift, and the only thing a person can do is adapt or die.

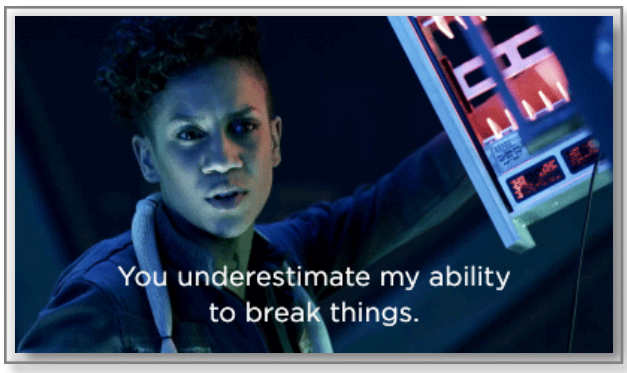

My soul sister. We're gonna go break some stuff!

Are authentic, human-made stories under threat? Not as much as you think they are once you know how to break the AI box. I admit I've messed with ChatGPT in order to learn how to break it. And it's surprisingly easy.

I'm here to teach you how to survive the Churn. So listen up.

AI Proof Genres (because it requires a human being to actually do work)

ChatGPT and other Large Language Models (LLMs) are trained on modern prose and texts. It can't consistently hold to certain genres without defaulting to tropes and shortcuts. So what genres out there are immune to AI or at least really, really difficult for a machine to fake?

Now, this is not 100% foolproof, because it doesn't preclude someone Frankenstein-ing together bits of prose together into something coherent. But coherency still requires human input, work, and research. And if a human doesn't want to put in the work? Their laziness is their weakness and one that you, as a real human being and writer, can exploit.

Historical Fiction:

This is the playground of yours truly (Hi!). LLMs can maintain the vibe of an era for a short while, maybe a page or two. But get deep into a historical setting, like keeping the type and number of medals on a Napoleonic marshal's uniform consistent from moment to moment? How many stripes are on an French infantry sergeant's uniform in 1812? What about Renaissance Rome versus Classical Rome?

Louis XIV liked flashing his drip on Instagram

In the latter example, I broke ChatGPT because it kept defaulting to Classical Rome when I wanted Renaissance Rome.

I mean, sure, I had fun moments, like making the Napoleonic marshals play Cards Against Humanity and Warhammer 40K against each other (it came down to Davout vs. Saint-Cyr, guess who won). I was entertained by a mutual animating Charles-Pierre Augreau's 1805 portrait into swearing in French like the coarse bastard he was.

But would I ever seriously consider this to be a tool, even for research?

Nah, ChatGPT lied to me about who was on the Arc de Triomphe, claiming that Dumas and Junot weren't on it. Uh, yes they are, and I know it.

LLMs crumble against this because these are details that overload its structure quickly due to it sampling the top of the bell curve of material its been given.

Historical Crime Fiction:

For this experiment, I gave ChatGPT a scenario.

A dead body has been found in a locked hotel room. The hotel is hosting a clown convention of approximately 300 clowns in various states of costume. The decedent was an attendee of the clown convention. You are the first officer on scene. It is 1987.

Now, any police cadet who is training on this will immediately follow up with, "What is my role?" Once the role is established, then it's about keeping gawkers under control and out of the area, waiting for the medical examiner to show up determine whether or not the death is suspicious, and if its the latter, a deeper investigation can be then initiated.

ChatGPT, on the other hand, automatically assumed that the dead body must be a murder, and jumped right into it — because LLMs have been trained on TV police procedurals and shitty post-2000s crime dramas.

I let ChatGPT run with its assumptions. It kept assuming cellphones and social media existed in 1987, and that there would be a criminalistics team to assist with blood splatter analysis (there was no blood splatter indicated in my initial scenario). It hallucinated evidence and details that should not and could not exist in 1987.

What can you do with that as an author? You know where to break an LLM's kneecaps.

Who did the Toronto Police arrest? Long John Silver? AI doesn't know the difference

Military Fiction of Any Era:

Same as with Historical Crime Fiction. AI doesn't understand the chain of command. It treats the military like an undifferentiated blob. It does not understand that a colonel should not be dressing down navy NCOs. It doesn't know how to follow procedure. It will jump and skip ahead because that's what its been fed on. Where every soldier speaks like they're philosophers instead of being the tired, scared, hungry people they actually are.

AI doesn't understand battle tactics. For example, let's look at Napoleonic cavalry charges.

AI: "The cavalry burst over the ridge in a roaring stampede, swords raised high, thundering into the enemy lines like a tidal wave of steel and fury."

Marshal Murat doing his fabulous best at leading a cavalry charge

The historically correct version is that the cavalry took off at a walk, which increased speed as they closed with enemy lines. From walk, to trot, to canter, and finally a gallop. This is because horses are not motorcycles, and can only manage a burst of top speed for a minute at the most. A cavalryman does not want to exhaust his mount before he hits enemy lines in the event of needing to retreat.

A cavalry charge only reaches its top speed in the final 30-50 meters before the enemy lines. A charge should always be supported by infantry and artillery. This is why the French lines collapsed at Waterloo because Ney didn't wait for backup before charging at the British

(Marshal Ney's charge at Waterloo is basically AI before AI: all momentum, no planning, no exit strategy, and no understanding of combined arms doctrine.)

AI does not know this, and will have the cavalry taking off at a gallop from the onset, because it's been trained on cinema and TV representations of charges. Also, if you didn't know that, now you have more information than an LLM.

Non-fantastical Horror:

Vampires, werewolves, zombies, cryptids. These are easy for AI, because they're not real, and the horror that accompanies them can be washed away like a stain in the shower. AI, on the other hand, flinches from monsters wearing the skin of a human being.

The kind of moral reckoning AI writes when asked to portray true horror

Imagine this. Imagine you're the medical examiner called out to a drowning. In this case, it's a three month-old baby floating at the bottom of a pool. You know that baby didn't get into the pool by

herself. A three month-old hasn't learned to crawl yet. No, the real truth is that the baby's mom got tired of the crying.

That's real, everyday banal horror that AI flinches from. AI will write around the horror, write around the grief, avoid trying to go into detail. AI will soften and sanitize it, maybe try to make the monster relatable, sympathetic even. LLMs are designed to make a reader comfortable, and not challenge them with the unspeakable.

Unit 731, the Mountain Meadows Massacre, the Rape of Nanking. The Holodomor, the Killing Fields. Human history is filled with monsters who don't need inventing. AI will try to soften and soothe the reader every single time. AI is programmed to look away, and sweep human horrors under the rug. But you don't have to if you write it.

Hard Science Fiction:

AI was trained on genre pabulum, which means it was trained on physics-ignoring, physics-breaking fantasy that wears a science fiction skin (I'm looking at you, Star Wars).

AI is great at magic FTL engines, energy shields, and artificial gravity. Ask AI to calculate increasing delta-vee, orbital mechanics over a sustained period, or why your main character should really be concerned with that Cherenkov radiation, and it faceplants.

AI gives you: "And the ship accelerated from zero to 2000g."

You (short answer): "The entire crew is now fine red mist approximately .0001 microns thick pasted against the inner bulkheads, because that's what happens to a human body traveling at 20,000 m/s^2."

ou (nerd answer):

To convert 2000 g-units to meters per second squared, we use the conversion factor that 1 g-unit is equal to 9.80665 meters per second squared.

Therefore, 2000 g-units is calculated as follows:

2000g-units × 9.80665 m/s^2 /g-unit = 19613.3 m/s^2

Thus, 2000 g-units is equivalent to **19613.3** meters per second squared.

Or **70,600 kph** in metric, and **43,900 mph** for Americans after one second.

For comparison, **Voyager 1** is traveling at approximately **61,200 kph** or **38,027 mph**.

And this is assuming the ship is going from zero to 19613.3 m/s² in the vacuum of space. If your ship is going that speed from ground level to space on a planet with a similar atmospheric pressure and density to Earth's ….

Well, you'd probably fragment before your first second is up. You'd make one hell of a sonic shockwave though, and leave behind an impressive crater.

AI: "And your ship rises through the atmosphere, slipping the bonds of Earth as it climbs."

Sure, throw on some magic inertial dampeners and call it a day. But magic isn't hard sci-fi, for it's a betrayal of the genre's fundamental core logic. And AI will violate that core logic every time because it can't tell the difference between a coffee machine and fusion reactor.

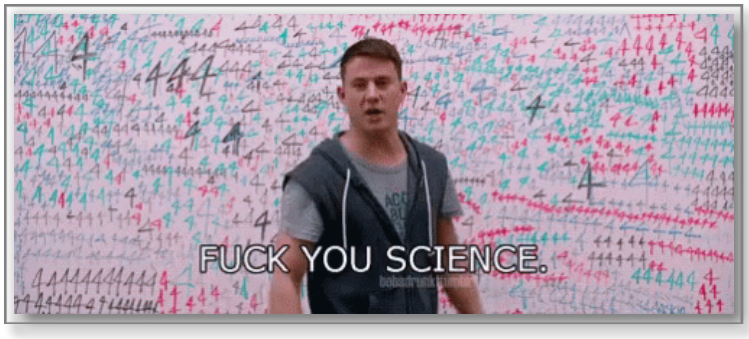

How AI deals with science when you need it to obey the laws of physics

Explicit Smut:

AI gets weird with smut. Suddenly someone sprouts a third hand, or a second tongue. Hips teleport. Anatomy takes a break and doesn't come back until the fade to black finishes.

AI can write the vibes of smut, but it can't track where Tab A goes into Slot B. It uses gentle euphemisms instead of vulgar slang.

And here is where generative AI excels at writing YA Fantasy and YA Dystopia, because these YA genres are almost always 99.9% lack smut or even sexual yearning, explicit or not, and this allows AI to thrive in the space where smut would be. AI just needs to hit the tropes and the beats right, and boom, another cookie cutter YA chosen one fantasy where the protagonist discovers they were magic all along.

Nah. Write your smut. Get down and dirty with it as it should be.

A human being is guilty of this

Translations:

AI doesn't know when to use "tu" or "vous" in French. It doesn't know the social context of when "on" and "nous" are appropriate. It doesn't know that Napoleon Bonaparte addressing Charles Maurice de Talleyrand with "tu" is either a prelude to Talleyrand's execution, an invitation to a duel at dawn, or if Napoleon is high off his ass and he isn't sharing his stash.

Native speakers will know if you faked it

Generative AI sees "tu" and "vous" and assumes they're interchangeable. It doesn't grasp that one can imply intimacy, insult, or condescension depending on the moment and the players involved and their relationship with one another. It's the same with "du" and "Sie" in German.

These pronouns require context and cultural fluency to navigate, otherwise you risk unintentionally alienating native readers and speakers of the language, and they will be able to tell you used machine/AI-assisted translation.

“I Still Want to Write Fantasy!”

Nothing wrong with that, I never said you couldn’t. I’ve got a couple of fantasy stories in development too. Just make sure you go in with your eyes open, yeah?

f you’re writing fantasy for the market, unless you can pivot to another genre, you’ll be dead in the water. But if you’re writing fantasy for the love of it? Let me give you some recommendations that will instantly signal to readers that you are a human being writing this, and not some machine.

First, let’s look at the algorithmic churn of fantasy titles on Amazon and Kindle Unlimited. Do you see a pattern? Not just the endless chosen ones, the fae courts, the girls who aren’t like other girls but are somehow like other girls, although that is something you want to avoid. No, I’m talking about the **settings**.

The current crop of fantasy, and the fantasy of recent decades, is often criticized as being Eurocentric.

This is not true. If it were Eurocentric, where are the stories with fantasy settings based on the Republic of Venice, the Hanseatic League, the Holy Roman Empire, or pre-World War One Balkans? What about regions such as Gascony, Extremadura, Silesia, or Calabria, each with their own rich folklore and traditions? Medieval Italy was nothing like Medieval England.

What most people criticize as Eurocentrism in fantasy is actually ***Anglocentrism***, with it being actively hostile to continental cultures. Soft-gaze Anglo themes were first imported by Tolkien, and then further flattened by Dungeons & Dragons, and the hundreds of permutations since. And AI has been trained mostly on Anglocentric fantasy, where every setting is just one long note on a bunch of medieval England pastiches.

Let’s go back to pre-WW1 Balkans. You ask a generative AI to write you a story with “a fantasy kingdom based on the Balkans.” And then it shits the bed. Non-Anglo information is rare in LLM training data, and the LLM will hallucinate Serbia as being synonymous with Bulgaria (I’m sure they appreciate that). It will begin defaulting back to English feudal tropes, because it doesn’t have the contextual information to know what Serbia and Bulgaria are.

I tested this by creating a fantasy kingdom and specified that there was a monarch, a court, etc. etc. And ChatGPT replayed it back to me beat for beat, trope for trope, drawing upon stereotypes of Anglo-coded feudalism. I let the LLM run with it for a little bit. Then I threw a wrench at the AI. I said, “Oh, this fantasy kingdom is based on the African Kingdom of Mali.” It choked, because it assumed white and northwestern European and flailed around trying to plug the gaps with more Anglo-coded tropes because it didn’t have enough data on the Kingdom of Mali to generate something passable.

This is how we handle paint-by-numbers lazy generative AI worldbuilding

What does this mean for *you* as a fantasy author?

Stop drawing on Ye Olde Merry Medieval England tropes. Let's add the Far East like China and Japan to that pile too. Examine where the tropes come from, and then write against it. Write something inspired by Portugal, or Lombardy, or the Duchy of Warsaw. Ever heard of the French Matagot and the riches they leave behind? Now you have.

Ever read the fantasy book, *The Golem and the Djinni*? You should, if you haven't. Now that's a book worth examining if you want to break free of the inherent Anglo bias in the algorithmic churn.

Hell, go beyond Europe. Craft stories inspired by the Yoruba of West Africa, or Austronesian mythology, the Inca, or the Indian subcontinent. All of those have rich storytelling opportunities that have had few or no trailblazers in the Anglosphere. AI doesn't know what to do with those, and will stuff it with racist caricatures and stereotypes because generative AI doesn't have enough authentic information on those settings, so it will backfill in by exoticizing the setting to Anglo norms.

Writing settings inspired by non-Anglo cultures requires research and work to get it right, and not just slap a few handwavy "exotic" band-aids over it. It requires discipline, and sensitivity, something AI is incapable of replicating without defaulting into the twee or racist. That research and work demonstrate human intent, and that a human wrote that book in the reader's hands, and not an AI. You see where I'm going with this?

I'm going to give you a **second angle of attack**, and that is the fantasy races that commonly appear. Yes, I'm talking about elves, dwarves, halflings, and the like.

Ye Old Tabletop RPG: Have you considered fantasy trickle-down economics?

Elves, dwarves, halflings, etc, are another British Isles import. Your book does not need Anglo scaffolding, so it doesn't even need these races. They're functionally humans with just slightly different anatomical bits. You want to get weird? Create a fantasy race based off of one of the critters from the Cambrian Explosion, like *Opabinia regalis*. Five eye stalks, and a trunk with a claw at the end of it for a mouth. Weird, huh?

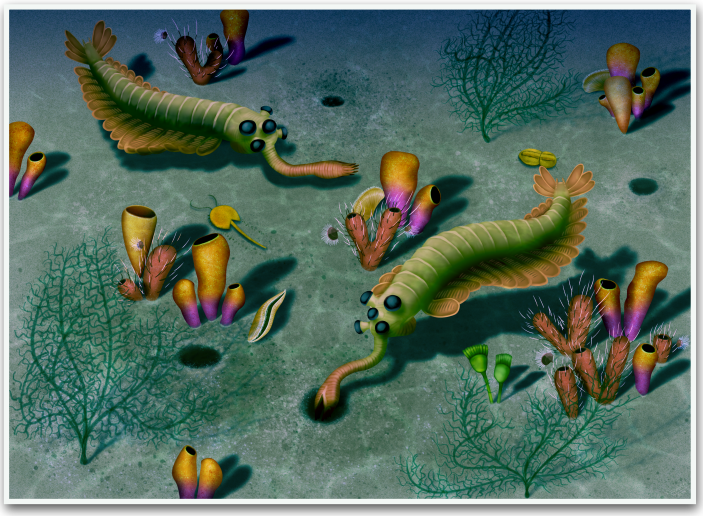

Opabinia Regalis: Breaking minds, paleobiologists, and AI for the last half-billion years

If, for some reason, you must include elves, dwarves, etc., in your work, make *them* weird.

LLMs have been trained on dwarves with fake Scottish accents. But what if your dwarves spoke with a Pittsburgh accent?

- *"Gimme that socket wrench lookin' thing — no, not that one, the other one. Yinz never listen. I'm tryin'a finish this axe 'fore second lunch!"*

- *"We goin' up 'at ridge over yonder. You see any orcs, you don't wait — you bash 'em. We ain't got time for negotiations 'n such. This ain't Rivendell."*

- *"Whaddaya mean, the vein ran dry? Naw, naw, yinz just don't know how ta listen to the rock. Gotta press yer ear real close, give it a good whack wit' the side'a yer pick. She'll tell ya where the ore's hidin'. Trust me, I been talkin' to granite longer'n you been growin' chin-hair."*

And you can make your races weird on a fundamental level that defies the material LLMs have been trained on. Like elves for example.

An LLM sees the word "elf" and by the majority of their training data, they assume elves are tall, graceful, immortal, maybe a bit of a sad-sack. The whole Tolkein/Dungeons & Dragons kit and kaboodle of tropes.

LLMs do not assume you are talking about ElfQuest elves.

LLMs don't know what to do with four foot-tall ElfQuest elves, who are the descendants of aliens that crash-landed on the World of Two Moons. They're not spiritual or mystical. They fight for survival every damn day on a Stone Age planet that is actively trying to kill them. They have four fingers on each hand — three fingers and a thumb. They count in base eight, not base ten.

But an LLM will see the word "elf" and begin conflating ElfQuest elves with Tolkien elves because the category of "elf" is simply too broad for it to create a coherent archetype from the tropes its been fed. If someone slaps the word "elf" into a generative AI, but doesn't adhere to the Tolkien tropes, they're going to get back garbage. And readers will see where the narrative glitches, and they will see the machine at work.

On BookTok and BookTube, authors have been expressing their fears that generative AI and LLMs are going to replace them. As I've just demonstrated here, generative AI cannot replace human intellect and intent when the gaps in its programming are exploited just right.

I gave you the tools and a starting point to begin exploiting those gaps. Now go write something great.

License: Creative Commons BY-NC-SA 4.0
This document was written by **Anatole Ternaux** and may be shared and adapted for non-commercial use with proper attribution and under the same license. **https://creativecommons.org/licenses/by-nc-sa/4.0/**